

TG-VQA: Ternary Game of Video Question Answering

Hao Li^{1,2,3*}, Peng Jin^{1,3*}, Zesen Cheng^{1,3}, Songyang Zhang², Kai Chen², Zhennan Wang⁴,
Chang Liu^{5†}, Jie Chen^{1,3,4†}

¹School of Electronic and Computer Engineering, Peking University, Shenzhen, China

²Shanghai AI Laboratory, Shanghai, China

³AI for Science (AI4S)-Preferred Program, Peking University Shenzhen Graduate School

⁴Peng Cheng Laboratory, Shenzhen, China

⁵Department of Automation and BNRist, Tsinghua University

lihao1984@pku.edu.cn, {jp21, cyanlaser}@stu.pku.edu.cn, sy.zhangbuaa@gmail.com,
chenkai@pjlab.org.cn, wangzhennan2017@email.szu.edu.cn, liuchang2022@tsinghua.edu.cn,
chenj@pcl.ac.cn

Abstract

Video question answering aims at answering a question about the video content by reasoning the alignment semantics within them. However, since relying heavily on human instructions, *i.e.*, annotations or priors, current contrastive learning-based VideoQA methods remains challenging to perform fine-grained visual-linguistic alignments. In this work, we innovatively resort to game theory, which can simulate complicated relationships among multiple players with specific interaction strategies, *e.g.*, video, question, and answer as ternary players, to achieve fine-grained alignment for VideoQA task. Specifically, we carefully design a VideoQA-specific interaction strategy to tailor the characteristics of VideoQA, which can mathematically generate the fine-grained visual-linguistic alignment label without label-intensive efforts. Our TG-VQA outperforms existing state-of-the-art by a large margin (more than 5%) on long-term and short-term VideoQA datasets, verifying its effectiveness and generalization ability. Thanks to the guidance of game-theoretic interaction, our model impressively convergences well on limited data (10^4 videos), surpassing most of those pre-trained on large-scale data (10^7 videos).

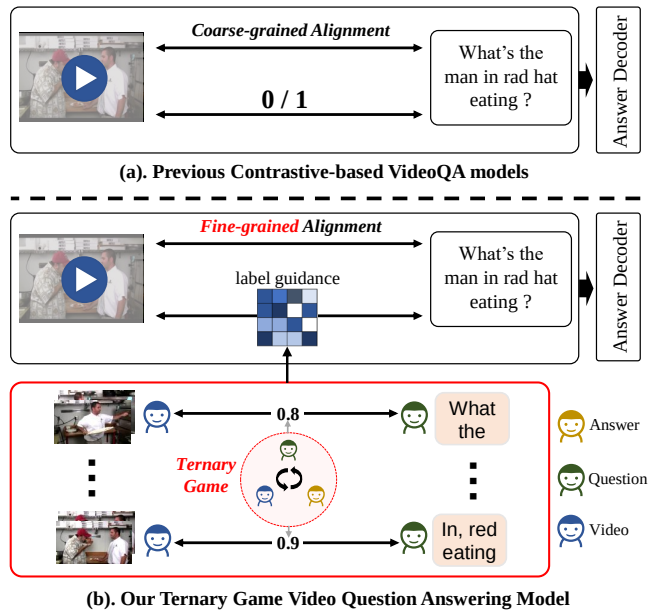


Figure 1: (a). Contrastive-based VideoQA models only learn a coarse-grained global alignment before the answer decoder. (b). To achieve fine-grained alignment, we model video, question, and answer as ternary game players and use a VideoQA-specific interaction to generate the label guidance for improvement.

1 Introduction

Video question answering (VideoQA) [Wu *et al.*, 2017; Sun *et al.*, 2021] aims to automatically infer the correct answer given a video and a related textual question. Such a multi-modal vision-language task has potential application in a broad range of applications such as vision language navigation in embodied AI [Gu *et al.*, 2022], video content retrieval by users’ questions [Jin *et al.*, 2022; 2023]. Tremendous progress has been made recently in VideoQA, thanks to advances in vision language pre-training and development in model architecture. However, for a VideoQA task, due to the intrinsic property of the visual data, we typically learn from

datasets with the *long sequence frames*, which consist of various visual appearances and rich motion information.

The visual data’s long sequence property introduces many challenges for multi-modal reasoning in the wild, as a deep learning model has to simultaneously cope with multi-modal representation learning, visual-linguistic alignment, and answer prediction [Li *et al.*, 2022e; 2022b]. The naive method would require high-quality annotated data and are typically data-hungry, and remains challenging to achieve accurate visual-linguistic alignment.

Early works on VideoQA focus on developing specific architectures [Jiang and Han, 2020; Li *et al.*, 2021; Qian *et*

et al., 2022] to align the visual representation and the linguistic question, which require sophisticated and heuristic design. More recent efforts [Lei *et al.*, 2021; Yang *et al.*, 2022a] aims to learn a VideoQA model with contrastive learning by leveraging the power of large-scale dataset, which rarely explore the fine-grained visual-linguistic alignment (shown in Figure 1(a)). This severely limits its modeling capacity and generalization ability in answer prediction. Thus, we raise a question: *can we achieve the accurate and robust fine-grained alignment in a data-efficient manner for VideoQA?*

To answer this question, we need to tackle the acquisition of the accurate annotation of fine-grained alignment between question semantics and video clips. However, it is prohibitive to collect the manual annotation due to the mega-scale of the video amount. A promising idea is to automatically generate the alignment annotation without labor-intensive effort.

Toward this goal, we focus on incorporating visual-linguistic alignment into the contrastive learning framework and propose an annotation generator of the fine-grained alignment (FAG) based on the *multi-player game theory* [Making, 2009]. Specifically, we introduce the interaction strategy for annotation generator construction, where the (*video*, *question*, *answer*) are treated as ternary game players, and the multi-player game theory could mathematically simulate the pairwise annotation between video clips and question semantics (illustrated in Figure 1(b)).

In this work, we first carefully design an interaction strategy to tailor the characteristics of VideoQA. Intuitively, if there exists a strong semantic correspondence between the *video player* and the *question player*, and these two players both have a large contribution to the answer, the coalition between them will be strengthened in our framework. Equipped with the annotation generator, we are able to model the fine-grained visual-linguistic alignment with additional supervision signals. To further improve the alignment efficiency, we also explore the multi-modal token reduction strategy for VideoQA. We thoroughly investigate different reduction methods and develop a clustering-based token merge module in the end. Our total framework is named Ternary Game VideoQA (TG-VQA).

We conduct extensive experiments to validate our model on three VideoQA datasets, including MSVD-QA, MSRVT-QA, and ActivityNet-QA. The empirical results and ablative studies show our method consistently achieves significant improvements (more than 5%) on all benchmarks. The annotation generator built from the ternary game also significantly improves the model convergence and data efficiency, which makes our TG-VQA competitive or superior compared with most pre-trained models learned from millions of video data. The main contributions are as follows:

- To the best of our knowledge, we are the first to bring game theory into VideoQA. Utilizing game theory’s ability to simulate the video-question token relations, our game theory-based annotation generator helps the VideoQA task achieve fine-grained alignment.
- Our alignment label generator is built from the ternary game. For the characteristics of the VideoQA task, the ternary Game models video, question, and answer

as ternary game players. The ternary game values the video-question pair’s alignment possibility and their contribution to the answer.

- We achieve new SoTA results on short-term and long-term VideoQA datasets, verifying the generalization ability. Without the pretraining stage, our TG-VQA also outperforms most VideoQA pre-trained models.

2 Related Works

2.1 Video Question Answering

The video question answering (VideoQA) task [Zhong *et al.*, 2022] requires models to analyze the complex semantic correlation between the video and the question. The VideoQA task has two main-stream models: (1). Hierarchical cross-attention-based models. (2). Contrastive learning-based models. Hierarchical cross-attention models [Xu *et al.*, 2017; Li *et al.*, 2019; 2022c; Peng *et al.*, 2022; Cai *et al.*, 2021; Fan *et al.*, 2019; Li *et al.*, 2023a] design Spatio-temporal attention structures to fusion the video and text features. Several recent models establish the effective alignment stage [Li *et al.*, 2021; Xiao *et al.*, 2021] for VideoQA. [Jiang and Han, 2020] constructs the video clips and text entities into heterogeneous graphs to achieve fine-grained alignment. [Li *et al.*, 2022e] establish the video-question alignment using invariant grounding. [Qian *et al.*, 2022] uses a locator to align the question with video segments. However, these alignment strategies cannot apply in contrastive learning frameworks due to their hierarchical attention structure.

Contrastive learning-based VideoQA models [Lei *et al.*, 2021; Kim *et al.*, 2020; Piergiovanni *et al.*, 2022] use contrastive loss for cross-modality explicit alignment and fusion. However, lacking fine-grained alignment annotations, they suffer from slow convergence, requiring massive video data [Bain *et al.*, 2021; Yang *et al.*, 2022a; Huang *et al.*, 2021; Li *et al.*, 2020] for pretraining. Therefore, we establish a ternary game-based contrastive learning VideoQA model, using game-theoretic interaction [Kita, 1999] to generate the fine-grained alignment annotations.

2.2 Game Theoretic Interaction

The game-theoretic interaction [Making, 2009; Ferguson, 2020] consists of a set of players with a revenue function. The revenue function maps each team of players to a real number which indicates the payoff obtained by all players working together to complete the task. The core of game-theoretic interaction is to allocate different payoffs to game individuals fairly and reasonably. There are several interaction strategies including core interaction [Jeukenne *et al.*, 1977], shapley interaction [Sun *et al.*, 2020] and banzhaf interaction [Marichal and Mathonet, 2011]. The game-theoretic interaction has multiple applications in different fields [Aflalo *et al.*, 2022; Datta *et al.*, 2016]. Recently, LOUPE [Li *et al.*, 2022d] uses two-player interaction as a vision-language pre-training task. In this paper, we design a new framework of ternary game interaction strategy for the VideoQA task.

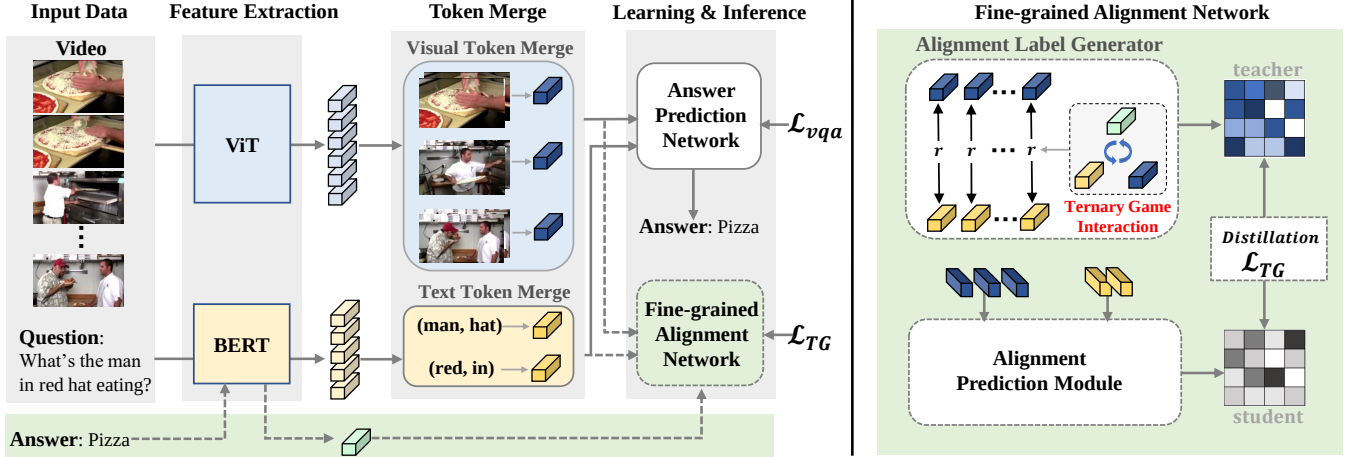


Figure 2: **The overall framework of our TG-VQA.** **Left)** We first use a *dual-stream transformer-based encoder* to extract the feature representation for the visual token and question token. We then introduce a *token merge network* to reduce the redundancy of the token and improve the efficiency for visual-linguistic alignment learning. Next, we use the *answer prediction network* to generate the answer for the input video-question pair. Moreover, we develop the *fine-grained alignment network* to explicitly align the visual token and question token at a fine-grained level. **Right)** The fine-grained alignment network consists of a *alignment label generator*, which is built from the ternary game interaction (video, question, answer ternary), and an alignment prediction module. We take the similarity matrix produced by the generator as the teacher and distill the fine-grained alignment knowledge from the ternary game interaction to the student. This extra supervision signals improve the consistency between the visual representation and linguistic representation and benefit the multi-modal reasoning in answer prediction. Notice that structures with **green background** and **gray dotted line** are auxiliary components, only used in the training process.

3 Preliminary of VideoQA and Game Theory

In this section, we first introduce the problem setting of video question answering in Sec.3.1, then briefly present the background for the multi-player game theory (Sec.3.2), which is utilized in our proposed alignment label generator.

3.1 Problem Setting of VideoQA

Given a video clip V and a text-form query Q , the VideoQA task aims to predict the correct answer \hat{a} from the answer space A . For the close-set type of question, A is a fix-size answering option list. For open-ended and multi-choice kinds of questions, A comprises the group of pre-defined answers and a list of candidate answering options. Generally, We formulate the VideoQA task as follows.

$$\hat{a} = \underset{a \in A}{\operatorname{argmax}} \mathcal{F}_\theta(a|Q, V) \quad (1)$$

where θ represents the trainable parameters group, \mathcal{F}_θ represents the modeling function.

3.2 Introduction of Game Theory

Toward the goal of achieving fine-grained alignment between video and question, we propose to leverage multi-player game theory to construct an alignment label generator. We aim to obtain the semantic relationship between visual tokens and question tokens, and their contribution to the answer. And the game theory targets generating an appropriate coalition construction strategy for multiple players. Thus, we propose to introduce game theory to align the label generation by considering the (video, question, and answer) as players.

Specifically, the multi-player game theory typically consists of (a) a set of players $\Gamma = (\mathcal{P}, \mathcal{R})$ consists a set of players $\mathcal{P} = \{1, 2, \dots, n\}$, and (b) a revenue function $\mathcal{R}(\mathcal{P})$. \mathcal{R}

maps each team of players to a real score, which indicates the payoff obtained by all players working together to complete the task. The key step of the game theory is to measure how much gain is obtained, and how to allocate the gain fairly.

In the multi-player game process, there are various different interaction strategies available, such as *Core interaction* [Jeukenne *et al.*, 1977], *Shapley interaction* [Sun *et al.*, 2020], *Banzhaf interaction* [Marichal and Mathonet, 2011]. Here we choose the Banzhaf interaction due to its balance of computational complexity and precision. Formally, given a coalition $\{i, j\} \in \mathcal{P}$, the Banzhaf interaction $\mathcal{B}(\{i, j\})$ for the player $\{i, j\}$ is defined as:

$$\mathcal{B}(\{i, j\}) = \sum_{C \in \mathcal{P} \setminus \{i, j\}} p(C) [\mathcal{R}(C \cup \{i, j\}) + \mathcal{R}(C) - \mathcal{R}(C \cup \{i\}) - \mathcal{R}(C \cup \{j\})], \quad (2)$$

where $\mathcal{P} \setminus \{i, j\}$ represents removing $\{i, j\}$ from \mathcal{P} , C stands for coalition. $p(C)$ is $\frac{1}{2^{n-2}}$, possibility for C being sampled.

Intuitively, $\mathcal{B}(\{i, j\})$ reflects the tendency of interactions inside $\{i, j\}$. The higher value of $\mathcal{B}(\{i, j\})$ indicates that player i and player j cooperate closely with each other. For VideoQA, we take the matrix \mathcal{B} as the alignment label annotation by changing the player definition and using a VideoQA-specific interaction strategy. We will start with a detailed description of our model architecture below.

4 Method

Our model consists of four main submodules: (1) a **backbone network** for generating feature representations of the video and question (Sec.4.1). (2) a **token merge network** for




Figure 3: We propose the token merge network to reduce the token redundancy. Specifically, we apply a 1D convolutional layer for temporal content encoding, a DPC-KNN for sparse token generation, and an attention layer for semantic fusion.

reducing the visual and question token number (Sec.4.1). (3) a **fine-grained alignment network** for establishing the fine-grained visual-linguistic alignment in VideoQA (Sec.4.2). (4) an VideoQA **answer prediction network** for generating the answer for VideoQA (Sec.4.3). We finally detail the training objectives and the inference pipeline in Sec.4.4. The overview of our proposed Ternary Game VideoQA (TG-VQA) model is illustrated in Figure2.

4.1 Backbone and Token Merge Network

Backbone. We adopt the ViT [Dosovitskiy *et al.*, 2020] and BERT as the backbone for generating visual representation and textural representation, respectively. Formally, we denote the representation of a video clip as a set of visual tokens $\mathbf{V} = \{\mathbf{v}_i | \mathbf{v}_i \in \mathbb{R}^{C_v}\}_{i=1}^{N_v}$, where \mathbf{v}_i is one frame feature vector with C_v channel and N_v is the total frame number. For the linguistic representation of question and answer, we first pad them into a fixed length N_l sequence and extract their textual feature via a transformer encoder initialized with BERT parameters. We formulate the generated question representation as a set of question tokens: $\mathbf{Q} = \{\mathbf{q}_j | \mathbf{q}_j \in \mathbb{R}^{C_l}\}_{j=1}^{N_l}$, where \mathbf{q}_j is a question token with C_l channel. Similarly, we are able to generate the representation of the answer \mathbf{A} with a text encoder, which is used to construct an alignment label generator.

Token Merge Network. To reduce the redundancy of visual tokens and question tokens, we develop a token merge network and investigate different merge strategies. Typically, the token merge network consists of (a) *temporal context encoding module*, (b) *sparse token generation module* and (c) *semantic fusion module*, which is illustrated in Figure 3. We will focus on the visual token in the remainder of this paragraph for notation clarity. Given a sequence of visual tokens \mathbf{V} , we first utilize the 1-D convolution layer to encode temporal context efficiently and denote the enhanced visual tokens as $\tilde{\mathbf{V}}$. Then, we conduct the sparse token generation for the $\tilde{\mathbf{V}}$ to reduce the number of tokens. Specifically, we investigate several different reduction strategies, like the random initialized sparse token(or called query) and clustering-based sparse token generation. We empirically find that the Density Peaks Clustering based KNN (DPC-KNN) [Rodriguez and Laio, 2014; Li *et al.*, 2023b] is superior for generating sparse and representative tokens, the ablative studies are shown in Sec.5.3.

we refer the reader to supplementary material for more details of the clustering method. Finally, we apply the cross-attention technique between the sparse tokens generated from the clustering technique and enhanced visual tokens $\tilde{\mathbf{V}}$, to further incorporate the semantic context information into the sparse tokens. We denote the spare visual tokens generated from the token merge network as $\mathbf{V}_s \in \mathbb{R}^{C_v \times N_{vs}}$, where N_{vs} is the number of sparse visual tokens and we have $N_{vs} < N_v$. Similarly, we are able to produce the spare textual tokens $\mathbf{Q}_s \in \mathbb{R}^{C_l \times N_{ls}}$ from the token merge network. Equipped with spare visual and question tokens, we not only reduce the redundancy in the input data but also improve the fine-grained visual-linguistic alignment efficiency (Sec. 4.2).

4.2 Fine-grained Alignment Network

Different from the current contrastive learning-based methods, which adopt the coarse-grained visual-linguistic alignment in model optimization, we develop a fine-grained alignment network that serves as an auxiliary network only existing in the training process, to explicitly supervise the model with an automatically generated alignment label.

The main idea is to introduce an alignment label generator to provide a supervision signal for learning a VideoQA model. We rethink the relationship among each item(video, question, and answer) and find that the alignment between visual tokens and question tokens actually reflects the semantic correspondence. Tokens sharing similar semantic meanings tend to simultaneously contribute to the final answer prediction. Thus, we propose to leverage the multi-player game theory to help find tokens with high semantic similarity. The fine-grained alignment network is composed of (a) an alignment label generator constructed with ternary game theory, and (b) an alignment prediction module.

Alignment Label Generator. Given the spare visual tokens \mathbf{V}_s and sparse question tokens \mathbf{Q}_s , we consider video, question and the answer \mathbf{A} as game players, which means $\mathcal{N} = \mathbf{V}_s \cup \mathbf{Q}_s \cup \mathbf{A}$. Intuitively, if a visual token has strong semantic correspondence with a question token, then they tend to cooperate with each other and contribute to the final answer. We now present the **ternary game interaction** strategy used in our work. For simplicity, we apply Banzhaf [Marichal and Mathonet, 2011] interaction for the ternary game. Concretely, a task-specific revenue function is required for interaction, and we need to consider the fine-grained visual-linguistic alignments as well as the token pair’s contribution to the final answer. Thus, the revenue function \mathcal{R} should satisfy the following criteria(we omit the subscript s of \mathbf{V}_s , \mathbf{Q}_s for clarity in the following paragraph):

- $\mathcal{R}(\mathbf{v}_i, \mathbf{q}_j)$ benefits from the semantic similarity between *video* and *question*.
- $\mathcal{R}(\mathbf{v}_i, \mathbf{q}_j)$ benefits from the semantic similarity between *the target answer representation A* and *the prediction from video-question pair*.

Thus, our proposed revenue function is formulated as:

$$\mathcal{R}(\mathbf{v}_i, \mathbf{q}_j, \mathbf{A}) = \phi(\mathbf{v}_i, \mathbf{q}_j) + \phi(\mathbf{A}, \mathbf{G}(\mathbf{v}_i, \mathbf{q}_j)). \quad (3)$$

#	Model	Initialization	Pretrain Data	MSRVTT-QA	MSVD-QA
Pretrained					
1	VideoCLIP [Luo <i>et al.</i> , 2022]	S3D+BERT	HowTo100M	33.8	31.8
2	ClipBERT [Lei <i>et al.</i> , 2021]	ResNet+BERT	HowTo100M	37.4	-
3	CoMVT [Seo <i>et al.</i> , 2021]	S3D+BERT	HowTo100M	39.5	42.6
4	VQA-T [Yang <i>et al.</i> , 2022a]	S3D+BERT	HowToVQA69M	41.5	46.3
5	ALPRO [Li <i>et al.</i> , 2022a]	ResNet+BERT	Web2M+CC3M	42.1	45.9
6	Co-Tok [Piergiovanni <i>et al.</i> , 2022]	K600+T5	How100M	45.7	48.6
7	FrozenBiLM [Yang <i>et al.</i> , 2022b]	CLIP+GPT3	WebVid10M	47.0	54.8
Non-Pretrained					
8	HCRN [Le <i>et al.</i> , 2020]	ResNet+LSTM	None	35.6	35.5
9	MHN [Peng <i>et al.</i> , 2022]	ResNet+LSTM	None	38.6	40.4
10	IGV [Li <i>et al.</i> , 2022e]	ResNet+BERT	None	38.3	40.8
11	VQA-T [Yang <i>et al.</i> , 2022a]	S3D+BERT	None	39.6	41.2
12	CLIP-QA [Radford <i>et al.</i> , 2021]	CLIP+BERT	None	39.0	38.5
13	CLIP4clip [Luo <i>et al.</i> , 2022]	CLIP+BERT	None	40.9	39.3
14	TG-VQA (ours)	S3D+BERT	None	42.7	45.5
15	TG-VQA (ours)	CLIP+BERT	None	46.3	52.5

Table 1: Experiments for the MSRVTT-QA and MSVD-QA datasets. We surpass the non-pretrained VideoQA models by a wide margin. Without the large-scale pretraining dataset, our interaction model also surpasses most of the pretrained VideoQA models.

#	Model	Initialization	Pretrain	Acc.
Pretrained				
1	VQA-T	S3D+BERT	69M	38.9
2	LF-VILA	Swin+BERT	8M	39.9
3	FrozenBiLM	CLIP+GPT3	10M	43.2
4	DeST	Swin+BERT	14M	46.8
Non-Pretrained				
5	LocAns	C3D+BERT	-	36.1
6	VQA-T	S3D+BERT	-	36.8
7	TG-VQA (ours)	CLIP+BERT	-	48.3

Table 2: Experiments of ActivityNet-QA (long-term). Our method surpasses all pretrained and non-pretrained VideoQA models.

where ϕ is a distance measurement for the semantic similarity, \mathbf{G} is a linear layer to project the concatenation of \mathbf{v}_i and \mathbf{q}_j into the answer representation space.

Then we apply \mathcal{R} to Eq.3.2 for the Banzhaf interaction. However, we find brute-force computation of Eq.3.2 is time-consuming. To speed up the interaction calculation process, we propose a deep learning-based idea by using a tiny convolutional network to predict the revenue function. Specifically, so we first calculate 1000 samples’ guidance matrix. We can use these matrices as data samples to learn a tiny model with short epochs. We take such a model as the alignment label generator to generate the guidance matrix (V, Q) .

Alignment Prediction Module. We adopt the contrastive learning-based framework to optimize the VideoQA model similar to [Lei *et al.*, 2021]. Differently, we introduce the explicit supervision signal for fine-grained visual-linguistic alignment. We first generate the alignment prediction $\mathbf{R}(\mathbf{V}_s, \mathbf{Q}_s)$ between visual tokens and linguistic tokens by computing their similarity. Then, we regard the guidance

matrix generated from the alignment label generator as the teacher and the alignment prediction as the student. The model is optimized by minimizing the Kullback-Leibler divergence between teacher and student. Ternary game loss is

$$\mathcal{L}_{TG} = \mathbb{E}_{\mathbf{V}_s, \mathbf{Q}_s} [\text{KL}(\mathcal{R}(\mathbf{V}_s, \mathbf{Q}_s), \mathbf{R}(\mathbf{V}_s, \mathbf{Q}_s))] \quad (4)$$

with such a distillation process, the model is expected to learn the multi-modal representation with rich semantic information and fine-grained visual-linguistic alignment.

4.3 Answer Prediction Network

Due to the established visual-linguistic fine-grained alignment benefit from the FAN, we are able to adopt a simplified answer prediction network, without the need for sophisticated multi-modal fusion/reasoning stages like many previous VideoQA models.

Specifically, given the sparse visual tokens \mathbf{V}_s and sparse textual tokens \mathbf{Q}_s generated from the token merge network, we first predict the token-level fusion weight by applying the non-linear projection (typically using linear layer + sigmoid function) on each token of \mathbf{V}_s and \mathbf{Q}_s . We then obtain the global representation for each modality by conducting weighted sum in \mathbf{V}_s and \mathbf{Q}_s , respectively. We denote the generated global feature vectors as \mathbf{v}_o and \mathbf{q}_o . Next, we concatenate those two vectors followed by an MLP to predict the answer *logits*. We use cross-entropy loss between *logits* and ground truth answer to supervise the whole framework.

4.4 Training and Inference

Combining the \mathcal{L}_{TG} from the ternary game module and the cross-entropy loss from the vqa prediction module, the overall loss of our model is the weighted sum of both parts:

$$\mathcal{L} = \mathcal{L}_{vqa} + \alpha \mathcal{L}_{TG}. \quad (5)$$

α is the trade-off hyper-parameter for the ternary game. As shown in Figure. 2, the structures with green backgrounds

Model	Acc.	Alignment	Strategy	Acc.	TM	Strategy	Acc.
Baseline	39.0	Coarse-grained	-	39.0	✗	-	44.5
Baseline+FAN	43.1	Fine-grained	Fully-connect	43.1	✓	Temporal	43.7
Baseline+FAN+TM	46.3		Ternary Game	44.5		DPC-KNN	46.3

Table 3: (I). Ablation for our fine-grained alignment network (FAN) and the token merge module (TM). Both modules have benefits. (II). The ablation for the alignment strategy. Both fine-grained alignment methods surpass the coarse-grained baseline. Our ternary game strategy outperforms the intuitively fully-connected strategy. (III). Ablation for the cluster strategy in the token merge module.




Figure 4: (a). The ablation study for hyperparameter α , the weight of \mathcal{L}_{TG} . Our TG-VQA performs best when $\alpha = 0.5$. (b). The ablation study for epoch performance between different VideoQA models. With our ternary game module, the model converges faster. (c). The performance analysis for the question category. Our ternary game module significantly improves “what” and “how” question correctness.

and Dotted lines are auxiliary components, which only appear in training. During the inference process, TG-VQA only activates the answer prediction network.

5 Experiments

5.1 Datasets

We select multiple VideoQA datasets to comprehensively evaluate the effectiveness of our method on different-length videos. Following the VQA-T [Yang *et al.*, 2022a] setting, we choose two short video datasets (MSVD-QA, MSRVT-QA) and one long video dataset (ActivityNet-QA) as our evaluation benchmarks. **MSVD-QA** [Xu *et al.*, 2017] comprises 1,970 short clips and 50,505 question-answer pairs. The clip’s average length is 10 seconds and the questions are divided into five question categories: what, who, how, when, and where. All of them are open-ended. **MSRVT-QA** [Xu *et al.*, 2017] comprises 10K videos and 243K question-answer pairs. The question types are similar to what is included in the MSVD-QA dataset. However, the scenario of the video is more complex, with a longer duration of 10-30 seconds. **ActivityNet-QA** [Yu *et al.*, 2019] is a Human annotated and large scale VideoQA dataset. The dataset consists of 58,000 QA pairs on 5,800 complex long web videos derived from the popular ActivityNet dataset. The average video length of ActivityNet-QA is 180 seconds, which is much longer than MSRVT-QA and MSVD-QA.

5.2 Experimental Results

We select the most recent pretrained and non-pretrained VideoQA models for comparison. Table. 1 shows the experimental results on the MSRVT-QA dataset and MSVD-QA

dataset. Compared with the non-pretrained VideoQA models, our method achieves sustainable improvements, 5.4% on MSRVT-QA and 11.3% on MSVD-QA. Without millions of video-text pretraining data, our method also surpasses most of the pretrained VideoQA models. Table. 2 shows the experimental results on the long-term VideoQA dataset, ActivityNet-QA. Our model achieves 48.3%, surpassing all pretrained and non-pretrained VideoQA models.

5.3 Ablation Studies

We first explore each module’s contribution to the overall TG-VQA performance. In Table.3 (I), Both our fine-grained alignment network (FAN) and token merge module (TM) benefit TG-VQA performance. More ablations are as follows.

Effectiveness of Alignment Strategies. Table.3(II) explores the benefits of various alignment strategies for the VideoQA task. Comparing line1 and line2, fine-grained alignment in VideoQA yields 4.1% benefits from coarse-grained alignment. For fine-grained alignment strategies, our ternary game alignment outperforms the intuitive alignment.

Effectiveness of Token Merge Module. Our token merge module clusters the video and question tokens aiming to reduce the token amount for later interaction. In Table. 3 (III), we make several attempts at the clustering strategies. “Temporal” represents clustering the tokens by temporal or sequence order. “DPC-KNN” represents our clustering strategy. Compared with line 1 using Banzhaf interaction with no merge module, the temporal cluster has a negative impact on the model performance due to the Lack of semantic correlation. Meanwhile, our DPC-KNN clustering strategy adaptively merges tokens under the guidance of semantic similarity, surpassing other clustering strategies.

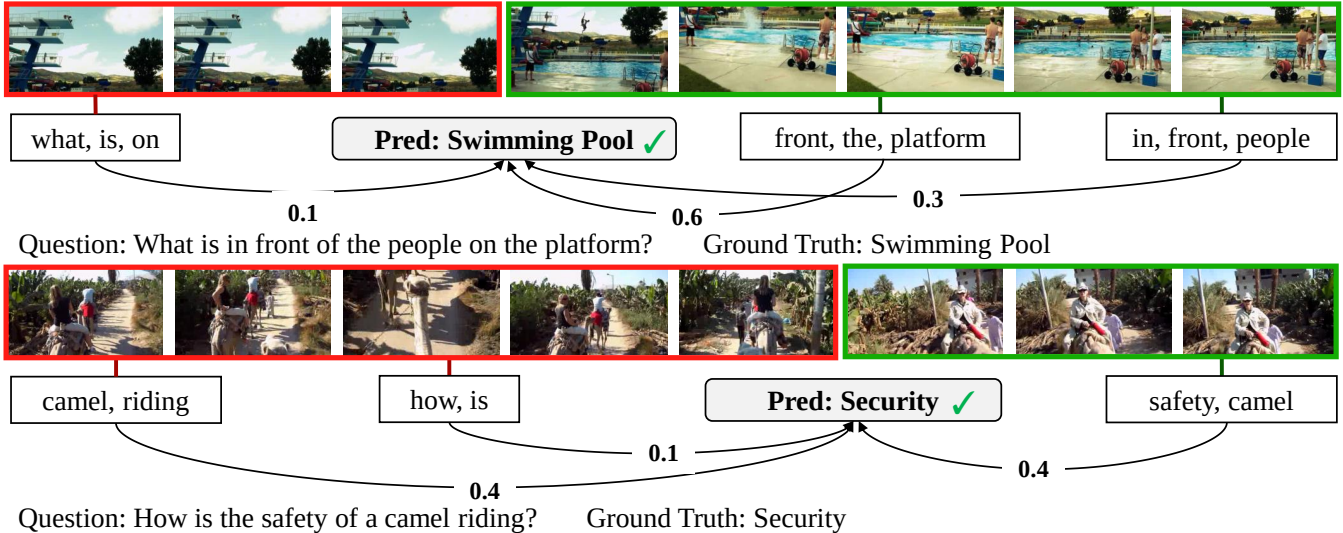


Figure 5: The case visualization. We visualize the most possible alignment pairs between video centers and question text centers. The arrow curves are the visualization of different video-question pairs’ contributions to the answer.

Impact of Encoders’ Initial Parameters. We list all VideoQA initial parameter combinations in Table. 1 row 3. Non-pretrained VideoQA models tend to use ResNet, S3D [Xie *et al.*, 2018] or CLIP as video encoder parameters while using BERT as text encoder parameters. Pre-trained models also adapt huge language models as the encoder, including GPT and T5 [Raffel *et al.*, 2020]. For a fair comparison, we apply the most common CLIP+BERT and S3D+BERT combinations. Shown in Table. 1, our TG-VQA model with S3D+BERT initialization surpasses other non-Pretrained VideoQA models with S3D+BERT. Using CLIP+BERT initialization, our TG-VQA model outperforms others by 5.4% in MSRVTT-QA and 11.7% in MSVD-QA. Due to FrozenBiLM’s large computation from GPT3 and WebVid10M, we don’t compare with FrozenBiLM.

Hyper Parameters in Train Objective. In order to explore the effect of the ternary game loss \mathcal{L}_{TG} ’s hyperparameter on the performance of the model, we train our TG-VQA on the MSRVTT-QA dataset with hyperparameter α from 0.1 to 1.5. Shown in Figure4 (a), the model performance fluctuates in range [45.1, 46.3]. When $\alpha = 0.5$, our model performs best.

Epoch Analysis. To illustrate our ternary game’s ability to accelerate the model’s convergence process with limited data, we visualize the epoch performance for our TG-VQA model and the CLIP4clip (non-pretrained VideoQA model without the ternary game) on MSRVTT. In Figure4 (b), both CLIP4clip and TG-VQA apply the same encoder initialization. With the fine-grained alignment network, our TG-VQA converges faster and better than CLIP4clip. We also visualize the epoch curve of the Co-Tok (a pretrained VideoQA model). With limited data, our TG-VQA surpasses Co-Tok on epoch 3, demonstrating our data efficiency.

Question Category Performance Analysis. We visualize the model’s performance on the Top-4 question categories on the MSRVTT-QA dataset. As shown in Figure4 (b), with the addition of our ternary game module, the model significantly improves “what” and “how” question types perfor-

mance, which attributes to the fine-grained alignment brought by the ternary game module.

5.4 Case Visualization

Figure 5 is the case visualization from the ActivityNet-QA dataset. For visualizing the cluster results, we cluster the video clips into two centers and the question tokens into three centers. Both cases show semantic similarity within the same centers. For visualizing the alignment results, Figure 5 shows the top-1 alignment pairs between the video center and the question center. The alignment results conform to the semantic consistency. For visualizing the contribution to the prediction answer, both cases illustrate that when a video-question pair is unlikely to be the answer, its contribution score is rather low (0.1). The second case illustrates that when multiple video-question pairs are similar to the answer, their contribution scores also tend to be the same. The visualization of two cases demonstrates the interpretability of our model.

6 Conclusion

In this paper, we study the fine-grained alignment in the VideoQA task. We innovatively model the VideoQA task as a ternary game process between video, question, and answer. We design a VideoQA-specific interaction strategy to simulate the alignment relationship. Experiments show the effectiveness, generalization, and data efficiency of our model.

Acknowledgements

This work was supported in part by the National Key R&D Program of China (No. 2022ZD0118201), Natural Science Foundation of China (No. 61972217, 32071459, 62176249, 62006133, 62271465), and the Natural Science Foundation of Guangdong Province in China (No. 2019B1515120049).

Contribution Statement

Hao Li and Peng Jin are Equal Contributions. Chang Liu and Jie Chen are Corresponding Authors.

References

- [Aflalo *et al.*, 2022] Estelle Aflalo, Meng Du, Shao-Yen Tseng, Yongfei Liu, Chenfei Wu, Nan Duan, and Vasudev Lal. VI-interpret: An interactive visualization tool for interpreting vision-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21406–21415, 2022.
- [Bain *et al.*, 2021] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021.
- [Cai *et al.*, 2021] Jiayin Cai, Chun Yuan, Cheng Shi, Lei Li, Yangyang Cheng, and Ying Shan. Feature augmented memory with global attention network for videoqa. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 998–1004, 2021.
- [Datta *et al.*, 2016] Anupam Datta, Shayak Sen, and Yair Zick. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*, pages 598–617. IEEE, 2016.
- [Dosovitskiy *et al.*, 2020] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [Fan *et al.*, 2019] Chenyou Fan, Xiaofan Zhang, Shu Zhang, Wensheng Wang, Chi Zhang, and Heng Huang. Heterogeneous memory enhanced multimodal attention model for video question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1999–2007, 2019.
- [Ferguson, 2020] Thomas S Ferguson. *A course in game theory*. World Scientific, 2020.
- [Gu *et al.*, 2022] Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason, and Xin Eric Wang. Vision-and-language navigation: A survey of tasks, methods, and future directions. *arXiv preprint arXiv:2203.12667*, 2022.
- [Huang *et al.*, 2021] Po-Yao Huang, Mandela Patrick, Junjie Hu, Graham Neubig, Florian Metze, and Alexander Hauptmann. Multilingual multimodal pre-training for zero-shot cross-lingual transfer of vision-language models. *arXiv preprint arXiv:2103.08849*, 2021.
- [Jeukenne *et al.*, 1977] J-P Jeukenne, A Lejeune, and C Mahaux. Optical-model potential in finite nuclei from reid’s hard core interaction. *Physical Review C*, 16(1):80, 1977.
- [Jiang and Han, 2020] Pin Jiang and Yahong Han. Reasoning with heterogeneous graph alignment for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [Jin *et al.*, 2022] Peng Jin, Jinfa Huang, Fenglin Liu, Xian Wu, Shen Ge, Guoli Song, David Clifton, and Jie Chen. Expectation-maximization contrastive learning for compact video-and-language representations. *Advances in Neural Information Processing Systems*, 2022.
- [Jin *et al.*, 2023] Peng Jin, Hao Li, Zesen Cheng, Kehan Li, Xiangyang Ji, Chang Liu, Li Yuan, and Jie Chen. Diffusionret: Generative text-video retrieval with diffusion model. *arXiv preprint arXiv:2303.09867*, 2023.
- [Kim *et al.*, 2020] Junyeong Kim, Minuk Ma, Trung Pham, Kyungsu Kim, and Chang D Yoo. Modality shifting attention network for multi-modal video question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10106–10115, 2020.
- [Kita, 1999] Hideyuki Kita. A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. *Transportation Research Part A: Policy and Practice*, 33(3-4):305–312, 1999.
- [Le *et al.*, 2020] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9972–9981, 2020.
- [Lei *et al.*, 2021] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [Li *et al.*, 2019] Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, Wenbing Huang, Xiangnan He, and Chuang Gan. Beyond rnns: Positional self-attention with co-attention for video question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8658–8665, 2019.
- [Li *et al.*, 2020] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. *arXiv preprint arXiv:2005.00200*, 2020.
- [Li *et al.*, 2021] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 34:9694–9705, 2021.
- [Li *et al.*, 2022a] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven CH Hoi. Align and prompt: Video-and-language pre-training with entity prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [Li *et al.*, 2022b] Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. Toward 3d spatial reasoning for human-like text-based visual question answering. *arXiv preprint arXiv:2209.10326*, 2022.
- [Li *et al.*, 2022c] Hao Li, Xu Li, Belhal Karimi, Jie Chen, and Mingming Sun. Joint learning of object graph and relation graph for visual question answering. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 01–06. IEEE, 2022.

- [Li *et al.*, 2022d] Juncheng Li, Xin He, Longhui Wei, Long Qian, Linchao Zhu, Lingxi Xie, Yueting Zhuang, Qi Tian, and Siliang Tang. Fine-grained semantically aligned vision-language pre-training. *arXiv preprint arXiv:2208.02515*, 2022.
- [Li *et al.*, 2022e] Yicong Li, Xiang Wang, Junbin Xiao, Wei Ji, and Tat-Seng Chua. Invariant grounding for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2928–2937, 2022.
- [Li *et al.*, 2023a] Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. Weakly-supervised 3d spatial reasoning for text-based visual question answering. *IEEE Transactions on Image Processing*, 2023.
- [Li *et al.*, 2023b] Kehan Li, Yian Zhao, Zhennan Wang, Zesen Cheng, Peng Jin, Xiangyang Ji, Li Yuan, Chang Liu, and Jie Chen. Multi-granularity interaction simulation for unsupervised interactive segmentation. *arXiv preprint arXiv:2303.13399*, 2023.
- [Luo *et al.*, 2022] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022.
- [Making, 2009] Making. Synthesis lectures on artificial intelligence and machine learning. *arXiv preprint*, 2009.
- [Marichal and Mathonet, 2011] Jean-Luc Marichal and Pierre Mathonet. Weighted banzhaf power and interaction indexes through weighted approximations of games. *European journal of operational research*, 2011.
- [Peng *et al.*, 2022] Min Peng, Chongyang Wang, Yuan Gao, Yu Shi, and Xiang-Dong Zhou. Multilevel hierarchical network with multiscale sampling for video question answering. *arXiv preprint arXiv:2205.04061*, 2022.
- [Piergiovanni *et al.*, 2022] AJ Piergiovanni, Kairo Morton, Weicheng Kuo, Michael S Ryoo, and Anelia Angelova. Video question answering with iterative video-text co-tokenization. In *European Conference on Computer Vision*, pages 76–94. Springer, 2022.
- [Qian *et al.*, 2022] Tianwen Qian, Ran Cui, Jingjing Chen, Pai Peng, Xiaowei Guo, and Yu-Gang Jiang. Locate before answering: Answer guided question localization for video question answering. *arXiv:2210.02081*, 2022.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [Raffel *et al.*, 2020] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
- [Rodriguez and Laio, 2014] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *science*, 344(6191):1492–1496, 2014.
- [Seo *et al.*, 2021] Paul Hongsuck Seo, Arsha Nagrani, and Cordelia Schmid. Look before you speak: Visually contextualized utterances. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16877–16887, 2021.
- [Sun *et al.*, 2020] Jianyuan Sun, Hui Yu, Guoqiang Zhong, Junyu Dong, Shu Zhang, and Hongchuan Yu. Random shapley forests: cooperative game-based random forests with consistency. *IEEE transactions on cybernetics*, 2020.
- [Sun *et al.*, 2021] Guanglu Sun, Lili Liang, Tianlin Li, Bo Yu, Meng Wu, and Bolun Zhang. Video question answering: a survey of models and datasets. *Mobile Networks and Applications*, 26(5):1904–1937, 2021.
- [Wu *et al.*, 2017] Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 163:21–40, 2017.
- [Xiao *et al.*, 2021] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9777–9786, 2021.
- [Xie *et al.*, 2018] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 305–321, 2018.
- [Xu *et al.*, 2017] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM Multimedia*, pages 1645–1653, 2017.
- [Yang *et al.*, 2022a] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Learning to answer visual questions from web videos. *arXiv preprint arXiv:2205.05019*, 2022.
- [Yang *et al.*, 2022b] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. *arXiv preprint arXiv:2206.08155*, 2022.
- [Yu *et al.*, 2019] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9127–9134, 2019.
- [Zhong *et al.*, 2022] Yaoyao Zhong, Wei Ji, Junbin Xiao, Yicong Li, Weihong Deng, and Tat-Seng Chua. Video question answering: Datasets, algorithms and challenges. *arXiv preprint arXiv:2203.01225*, 2022.